920584-906019

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In the application of	:	Clive C. Hayball
Serial No.	:	09/747,698
Filed	:	December 22, 2000
For	:	Network Proxy Apparatus and Methods
Examiner	:	Halim, Sahera
Art Unit	:	2157
Customer number	:	23644
Confirmation No.	:	9931

APPEAL BRIEF

Honorable Director of Patents and Trademarks
PO Box 1450
Alexandria, VA 22313-1450

Dear Sir:

This appeal is from the Examiner's final Office Action of February 6, 2008, confirmed in the Advisory Action mailed May 2, 2008 in which all pending claims (namely Claims 1-6, 8 and 13) were rejected. A timely Notice of Appeal was filed June 24, 2008 with the required fee.

This brief is being filed along with the required $510 fee pursuant to 37 C. F. R. § 41.20(b)(2).

(i) Real Party in Interest

This application is assigned to Nortel Networks Limited, who is the real party in interest.

(ii) Related Appeals and Interferences

There are no related appeals or interferences.

(iii) Status of Claims

This application was filed with claims 1 to 14. Claims 1-6, 8, 12 and 13 were elected in response to the election request dated November 16, 2004. Claim 12 was cancelled in the response to the Office Action dated March 23, 2005. In the responses of March 23, 2005, October 18, 2005, June 16, 2006, December 19, 2006, July 2, 2007 and February 6, 2008, claims 1-6, 8 and 13 were each amended at least once. It is the rejection of claims 1-6, 8 and 13 that is appealed, and those claims are set forth in the Claims Appendix.

(iv) Status of Amendments

A response was filed April 24, 2008, which has been entered, but the rejections were maintained in the Advisory Action of March 2, 2008.

(v) Summary of Claimed Subject Matter

Claim 1 recites a method of indexing location of content cached within an IP-based network comprising. The method includes the step of intercepting data traffic flowing from a source node to a destination node in the network as described on page 7 lines 20 to 23 and illustrated in Figure 2. The data traffic includes content to be cached at the destination node as discussed in page 7 lines 20 to 26 and 28 to 32.

The method also includes the steps of extracting identity information for the content and associated destination location information for the destination node where the content in the data traffic is to be cached from the data traffic (described on page 7

lines 25 to 29 with reference to the identity extractor and location extractor illustrated in Figure 3).

A mapping is generated from the content identified by the extracted identity information to the destination node identified by the associated destination location information as described on page 7 lines 31 to 32.

The mapping is stored in a content index database (page 8 line 1). The content index database is operable to provide an instance mapping containing a list of destination nodes at which the content has been cached as described on page 9 lines 2 to 4. The instance mapping is provided in response to an instance request containing the identity information for the content. (Page 8 Lines 17 to 18 and Figure 6).

Claim 2 references Claim 1 and recites that the step of intercepting data traffic is carried out by intercepting data traffic flowing into a cache as described on page 5 lines 8 to 9. Claim 2 further includes the further step of advertising the content identities for which mappings are stored in the content index by sending advertising messages to a predetermined location in the network (page 10 lines 1 to 2).

Claim 3 refers back to Claim 2 and includes the additional feature of recording the time of data traffic flows into the cache which are related to a particular content item and calculating the time period between a first flow of the content item into the cache and a subsequent flow of the content item into the cache thereby to assess how long items are held in the cache before they are expired (Page 12 lines 20 to 24). This allows the mapping relating to that content item to be deleted when that content item is judged to have expired in the cache as discussed on Page 8 line 33 to page 9 line 2.

Claim 4 references Claim 1 and further states that the step of intercepting data traffic is carried out by intercepting data traffic flowing out of an original content source node as described on page 7 lines 22 to 34.

Claim 5 refers back to Claim 4 and includes the additional step of receiving an advertising message, which advertises a mapping generated elsewhere on the network and which is related to content items stored in the original content source node, and augmenting the content index using information contained in the advertising message (page 5 lines 14 to 16 and Figure 7).

Claim 6 refers back to Claim 1 and includes the additional feature of "the step of intercepting data traffic... [being]... carried out by intercepting content requests issued by a cache". This is described, for example, on page 3 lines 28 to 29.

Claim 6 also includes the additional step of "advertising the content identities for which mappings are stored in the content index by sending advertising messages to a predetermined location in the network" (page 5 lines 9 to 12).

Claim 8 and 13 recite a proxy for an IP-based network and a computer program respectively, each being able to carry out the method of Claim 1, and the discussion of claim 1 above also applies to these claims.

(vi) Grounds of Rejection To Be Reviewed on Appeal

There are 2 rejections at issue:

1. the rejection of claims 1, 4, 8 and 13 under 35 USC § 103(a) as being unpatentable over Colby (US Publication 2005/0193114) in view of Tucker (US Publication 2004/0049598);

2. the rejection of claims 2, 3, 5 and 6 under 35 USC § 103(a) as being unpatentable over Colby (US Publication 2005/0193114) in view of Tucker (US Publication 2004/0049598), further in view of Rochberger (US Patent No. 6,205,146).

(vii) Argument

Ground 1

The Examiner has rejected Claims 1, 4, 8 and 13 as being unpatentable over US Pat. Pub. No. 2005/0193114 to Colby et al. (hereinafter Colby) in view of US Pat. Pub. No. 2004/0049598 to Tucker et al. (hereinafter Tucker).

Colby describes a "content-aware flow switch" that is linked to web servers and includes a contents server database (CSD) (Figure 2 and paragraph 50). The database includes "an associated IP address, URL, protocol, layer 4 port number, QoS indicators and the load balance algorithm to use" for each particular content (paragraph 51). Using this database the Content-Aware Flow Switch can select a server to provide content to a device that has requested it.

The database is constructed using ICP (Intelligent Content Probe). Colby states that "periodically, the ICP probes the servers front-ended by the content-aware flow switch for information regarding server status and content. This activity may be undertaken proactively (such as polling for general server health) or at the request of the CSD" (paragraph 59).

Additionally "the IPP [Internet Probe Protocol] periodically sends local server load and content information to neighboring content-aware flow switches. Data arriving from these peers is evaluated and appropriate updates are sent to the CSD" (paragraph 60).

Thus, the database in Colby is built up by specifically probing the web servers and neighbouring content-aware flow switches. Therefore, Applicant submits, and the Examiner agrees on page 4 of the Office Action dated February 6, 2008, that Colby does not disclose the features recited in Claim 1 of:

"(a) intercepting data traffic flowing from a source node to a destination node in the network, the data traffic including content to be cached at the destination node,

(b) extracting identity information for the content and associated destination location information for the destination node where the content in the data traffic is to be cached from the data traffic,

generating a mapping from the content identified by the extracted identity information to the destination node identified by the associated destination location information"

The Examiner, instead, contends that one skilled in the art would combine the disclosures of Colby and Tucker to arrive at the claimed invention. Applicant respectfully disagrees.

Tucker describes a network including distribution servers and control servers. When a request for content is received "the distribution server tries to fill the request from the cache manager, and if the content is present and current, the compressed or uncompressed content is provided to the user; otherwise, the distribution server tries to get the content from a control server. The control server also checks for the requested content in the cache and if the content is available and current, it is served to the user through the distribution server" (paragraph 32).

All that Tucker discloses is that "when the content is served to the user through the distribution server, it is compressed and cached in the distribution server" (paragraph 33).

Thus, Applicant submits that one skilled in the art would learn from Tucker to cache content that is being sent to an endpoint at a node. One skilled in the art would not

learn to extract associated destination location information for where the content is to be cached as this is not required if the content is to be cached at that node.

Applicant therefore submits that one skilled in the art would only learn to cache content flowing through a designated node and to maintain a list of content cached at the node.

Applicant therefore submits that neither Colby nor Tucker disclose or even suggest that the content should be intercepted at the node, identity information for the content and associated destination location information extracted, generating a mapping from the content identified by the extracted identity information to the destination location information as recited in Claim 1.

Furthermore, Applicant submits that one skilled in the art would not consider combining the features of Colby and Tucker. Colby and Tucker are both directed towards a similar problem, namely that of identifying and providing content within a distributed network. They both solve this problem in different ways, Colby by maintaining a database on a switch that is regularly updated and Tucker by creating extra caches of content at a switch.

Applicant submits that one skilled in the art, upon reading Colby and Tucker, would realize that they are two separate solutions and would have no reason to combine the teachings.

Nevertheless, even if one skilled in the art did combine the teachings of Colby and Tucker, which is not admitted, Applicant submits that they would merely learn to cache content at the Content-Aware Flow Switch in Colby in order that the web servers could be bypassed. They would not learn to create a mapping from content identity information and destination location information indicating where the content is to be cached.

For these reasons Applicant submits that Claim 1 is patentable over Colby in view of Tucker.

Applicant submits that Claims 8 and 13 which recite corresponding features to Claim 1 are also patentable over Colby in view of Tucker for at least the reasons given with reference to Claim 1.

Applicant submits that Claim 4 is patentable at least by virtue of its dependency upon Claim 1.

Ground 2

Applicant submits that Rocheberger also does not teach the features of

"(a) intercepting data traffic flowing from a source node to a destination node in the network, the data traffic including content to be cached at the destination node,

extracting identity information for the content and associated destination location information for the destination node where the content in the data traffic is to be cached from the data traffic,

generating a mapping from the content identified by the extracted identity information to the destination node identified by the associated destination location information"

Indeed, Rocheberger is not concerned with identifying where content is cached in a network. Rather, Rocheberger is concerned with identifying efficient routes to well known addresses. Well known addresses are registered with ports to enable the routing to be effected.

Therefore, Applicant submits that Claims 2, 3, 5 and 6 are patentable over the disclosures of Colby, in view of Tucker and in view of Rocheberger.

CONCLUSION

Given the above, it is submitted that the Examiner's rejections are in error, and reversal is urged.

August 22, 2008

Respectfully submitted,

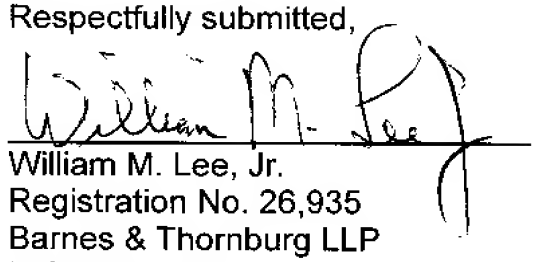

William M. Lee, Jr.
Registration No. 26,935
Barnes & Thornburg LLP
P.O. Box 2786
Chicago, Illinois 60690-2786
(312) 214-4800
(312) 759-5646 (fax)

Claims Appendix

1. A method of indexing location of content cached within an IP-based network comprising:-

   (a) intercepting data traffic flowing from a source node to a destination node in the network, the data traffic including content to be cached at the destination node,

   (b) extracting identity information for the content and associated destination location information for the destination node where the content in the data traffic is to be cached from the data traffic,

   (c) generating a mapping from the content identified by the extracted identity information to the destination node identified by the associated destination location information, and

   (d) storing the mapping in a content index database which is operable to provide an instance mapping containing a list of destination nodes at which the content has been cached, the instance mapping being provided in response to an instance request containing the identity information for the content.

2. A method according to claim 1, wherein the step of intercepting data traffic is carried out by intercepting data traffic flowing into a cache, and wherein the method further comprises advertising the content identities for which mappings are stored in the content index by sending advertising messages to a predetermined location in the network.

3. A method according to claim 2, wherein the method further comprises recording the time of data traffic flows into the cache which are related to a particular content item and calculating the time period between a first flow of the content item into the cache and a subsequent flow of the content item into the cache thereby to assess how long items are held in the cache before they

are expired and deleting the mapping relating to that content item when that content item is judged to have expired in the cache.

4. A method according to claim 1, wherein the step of intercepting data traffic is carried out by intercepting data traffic flowing out of an original content source node.

5. A method according to claim 4, wherein the method further comprises receiving an advertising message which advertises a mapping generated elsewhere on the network and which is related to content items stored in the original content source node, and augmenting the content index using information contained in the advertising message.

6. A method according to claim 1 wherein the step of intercepting data traffic is carried out by intercepting content requests issued by a cache, and wherein the method further comprises advertising the content identities for which mappings are stored in the content index by sending advertising messages to a predetermined location in the network.

7. (Withdrawn) A method of retrieving content in an IP-based network comprising the steps of:-

(a) intercepting a content request containing information related to the identity of a content item and a specified source location for the content item,

(b) sending an instance request to a content index associated with the specified source location, the instance request including the identity of the requested content,

(c) receiving an instance mapping from the content index which contains a list of instances and associated locations for the requested content,

(d) selecting the best instance of the content from the list,

(e) obtaining the requested content from the location associated with the best instance of the requested content, and

(f) returning the requested content to the requester of the content.

8. A proxy for an IP-based network comprising:-

(a) a data input operable to receive data transmitted from a source node to a destination node from the network, the data including content to be cached at the destination location,

(b) a data output operable to send data including the content to the network,

(c) an identity extractor operable to analyse data received at the data input and to extract, from the data, identity information for the content,

(d) a location extractor operable to analyse data received at the data input and to extract, from the data, location information for the destination node where the content in the data is to be cached,

(e) a mapping generator operable to generate a mapping from a content item identified by identity information provided by the identity extractor, to at least one destination node where the content is to be cached, the at least one destination node identified by associated destination location information provided by the location extractor, and

(f) a content index database operable to store a mapping provided by the mapping generator and which is operable to provide an instance mapping containing a list of destination nodes at which the content has been cached, the instance mapping being provided in response to an instance request containing an identity of the content item.

9. (Withdrawn) A proxy for an IP-based network comprising:-

(a) a data input operable to receive data from the network,

(b) a data output operable to send data to the network,

(c) a location requester operable to identify a request for a content item in data received at the data input and to send an instance request to a

content index associated with the source location of the content item specified in the content request, the instance request including the identity of the requested content, and

(d) a content returner operable to receive an instance mapping from the content index which contains a list of instances and associated locations for the requested content, to select the best instance of the content from the list, to obtain the requested content from the location associated with the best instance of the requested content, and to return the requested content to the requester of the content.

10. (Withdrawn) An advertising message for transmission over an IP-based network, the message being arranged to advertise a replica content item and a location for that item.

11. (Withdrawn) An instance request for transmission over an IP-based network, the request including a destination address and the identity of a requested content item, the destination address being a different address to the address of the source location of the content item.

12. (Cancelled)

13. A computer program, stored on computer readable medium, which, when executed by an apparatus causes the apparatus to index content in an IP-based network by:-

(a) intercepting data traffic flowing from a source node to a destination node in the network, the data traffic including content to be cached at the destination node,

(b) extracting identity information for the content and associated destination location information for the destination node where the content in the data traffic is to be cached from the data traffic flow,

(c) generating a mapping from the content identified by the extracted identity information to the destination node identified by the associated destination location information, and

(d) storing the mapping in a content index database which is operable to provide an instance mapping containing a list of destination nodes at which the content has been cached, the instance mapping being provided in response to an instance request containing the identity of the content.

14. (Withdrawn) A computer program which when executed retrieves content in an IP-based network by:-

(a) intercepting a content request containing information related to the identity of a content item and a specified source location for the content item,

(b) sending an instance request to a content index associated with the specified source location, the instance request including the identity of the requested content,

(c) receiving an instance mapping containing a list of instances and associated locations for the requested content,

(d) selecting the best instance of the content from the list,

(e) obtaining the requested content from the location associated with the best instance of the requested content, and

(f) returning the requested content to the requester of the content.

Evidence Appendix

None.

Related Proceedings Appendix

None.